EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	0	"10744444"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:16
L2	2	("6959390").PN.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:22
L3	1070	(identical$2 with (structur$3 or tree or hierarch$3) with storage$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 15:42
L4	64	((creat$3 or generat$3) with identical$2 with (structur$3 or tree or hierarch$3) with storage$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:24
L5	8774	((root or tree or hierarch$3) with key)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:43
L6	3	4 and 5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:25
L7	17	3 and 5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:25
L8	6	7 and @ad<"20001115"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:44

L9	61	(dual with key with (tree or hierarch$3 or structur$3)) and @ad<"20001115"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:29
L10	0	9 and 3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:30
L11	7	9 and 5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:30
L12	17549	(public with key) and (private with key)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 13:52
L13	1480	12 and 5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:43
L14	180763	((structur$3 or tree or hierarch$3) with storage$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:44
L15	335	13 and 14	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:44
L16	73	15 and @ad<"20001115"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:48

L17	123703	((structur$3 or tree or hierarch$3) with identical$2)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:45
L18	5	16 and 17	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:45
L19	1743	((structur$3 or tree or hierarch$3 or root) with storage$1 with key)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:47
L20	35	19 and 3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:47
L21	15	20 and @ad<"20001115"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 13:51
L22	66	19 and 12 and @ad<"20001115"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 12:52
L23	12	22 and (usage with key)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 13:53
L24	44	((migrat$5 or transfer$3) with storage with (tree or root or hierarch$3 or structur$3)) and (usage with key)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 15:28

EAST Search History

L25	81	((migrat$5 or transfer$3) with storage with (tree or root or hierarch$3 or structur$3)) and 3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 14:49
L26	212214	"225" and @ad<"20001115"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 13:52
L27	45	25 and @ad<"20001115"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 14:51
L28	0	27 and 12	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 13:52
L29	0	27 and (usage with key)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 13:53
L30	2	27 and (key with load$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 13:54
L31	0	27 and (key with (tree or hiearch$3 or root))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 14:51
L32	0	27 and 707/100.ccls.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 14:47

EAST Search History

L33	1	("7242768").PN.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 14:49
L34	1	"6212635".PN.	USPAT; USOCR	OR	OFF	2007/07/15 14:48
L35	1	25 and (securit$3 with key)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 15:29
L36	103	((migrat$5 or transfer$3) with storage with (tree or root or hierarch$3 or structur$3)) and (securit$3 with key)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 14:49
L37	22	36 and (key with (tree or hiearch$3 or root))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 14:51
L38	1	37 and @ad<"20001115"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 14:51
L39	30	36 and @ad<"20001115"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 15:42
L40	2	"016900".apn.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 14:53
L41	2	"016700".apn.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 14:53

C:\Documents and Settings\cnguyen14\My Documents\EAST\Workspaces\identical storage key.wsp

L42	3656	((migrat$5 or transfer$3) with (sign$3 or signature) with key)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 15:29
L43	632	42 and (securit$3 with key)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 15:29
L44	1	43 and 3	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 15:30
L45	184	43 and 5	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 15:30
L46	7	45 and (380/277).ccls.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 15:30
L47	41	((structur$3 or tree or hierarch$3) with storage$1 with key) and 43	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 15:42
L48	10	47 and @ad<"20001115"	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2007/07/15 15:42

Google

dual key tree	Search	Advanced Search
		Preferences

Web   Results **1 - 10** of about **1,920,000** for **dual key tree**. **(0.23 seconds)**

BROADCAST ENCRYPTION WITH DUAL TREE SIZES - Patent 20060109985
The method of claim 9 wherein an initial **key tree** is the sub-**tree** of an ..... [0031] While the
particular BROADCAST ENCRYPTION WITH **DUAL** SIZE TREES as ...
www.freepatentsonline.com/20060109985.html - 33k - Cached - Similar pages

Trusted computing platform with dual key trees to support multiple ...
Trusted computing platform with **dual key** trees to support multiple ... The second storage
tree (SST) will have a storage **key** for every migratable storage ...
www.freepatentsonline.com/20020059286.html - 39k - Cached - Similar pages

Broadcast encryption with dual tree sizes invention
The extended **key tree** may be designed such that the initial **key tree** is the ... Receive info
on patent apps like Broadcast encryption with **dual tree** sizes ...
www.freshpatents.com/Broadcast-encryption-with-**dual-tree**-sizes-
dt20060525ptan20060109985.php - 32k - Cached - Similar pages

[PDF] A Novel Dual-Key Management Protocol Based on a Hierarchical ...
File Format: PDF/Adobe Acrobat
cope with **key** management, such as **Tree Key** Graph (TKG) ... and the clustering
techniques, in this paper, we propose a **dual-key** management ...
www.springerlink.com/index/BL38JVYGVD27L9XC.pdf - Similar pages

PKI in the Enterprise - Course 586
Archiving keys automatically & manually; Accessing data with **key** or data recovery; Solving
deniability with **dual-key** support. Establishing PKI Policies ...
www.learning**tree**.com/courses/586.htm - 33k - Cached - Similar pages

Amazon.com: Secure Group Communications Over Data Networks: Books ...
blinded keys, node secrets, binary **key tree**, rekeying messages, new cluster **key**, **dual**
encryption protocol, collusion scenario, current group **key**, ...
www.amazon.com/Secure-Group-Communications-Over-Networks/dp/0387229701 - 129k -
Cached - Similar pages

[PDF] Evaluation of Dual-Structure Key-management Scheme suitable for ...
File Format: PDF/Adobe Acrobat
servers in the **dual**-structure scheme are lower than those in. the **tree**-based ..... the **key**-
management **tree** and h is the hight of the **tree**. f is. given by ...
doi.ieeecomputersociety.org/10.1109/MDM.2006.83 - Similar pages

Small Tree Offers Turn Key Solution For Mac. Xilinx Releases ...
SMALL **TREE** OFFERS HIGH PERFORMANCE TURN **KEY** SOLUTION FOR MAC ...
Available for $11499, this system includes: **dual** 2.5 GHz Power Mac G5; 8 GB RAM; ...
www.hpcwire.com/hpc/334196.html - 8k - Cached - Similar pages

[PDF] A Balanced Key Tree Approach for Dynamic Secure Group Communication
File Format: PDF/Adobe Acrobat
as the topology of its **key tree**. Both the AVL **tree** and the ..... [12] L. R. Dondeti, S.
Mukherjee, and A. Samal, "A **dual** encryption ...
ieeexplore.ieee.org/iel5/10216/32582/01523882.pdf?arnumber=1523882 - Similar pages

Evaluation of **Dual**-Structure **Key**-management Scheme Suitable for ...
Evaluation of **Dual**-Structure **Key**-management Scheme Suitable for Mobile Services ...
Tree-based schemes aim at reducing the load on the server, ...
portal.acm.org/citation.cfm?id=1136643.1136705 - Similar pages

1 2 3 4 5 6 7 8 9 10   **Next**

Try Google Desktop: search your computer as easily as you search the web.

| dual key tree | Search |

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

©2007 Google - Google Home - Advertising Programs - Business Solutions - About Google

Google

identical storage tree		Search	Advanced Search
			Preferences

Web   Results **1 - 10** of about **1,670,000** for <u>identical **storage tree**</u>. **(0.11 seconds)**

[PDF] <u>Variable-length Constrained-**storage Tree**-structured Vector ...</u>
File Format: PDF/Adobe Acrobat
wise design of constrained **storage** balanced **tree**-structured .....
distortion, the resulting distortions are **identical**, allowing a ...
ieeexplore.ieee.org/iel4/83/16166/00748888.pdf?arnumber=748888 -
<u>Similar pages</u>

> [PDF] <u>SSDBM'00: TSA-**tree**: A Wavelet-Based Approach to Improve the ...</u>
> File Format: PDF/Adobe Acrobat
> In this section, we formally prove that the **storage** space.
> required to store the optimal subset of TSA-**tree** (OTSA-. **tree**) is
> **identical** to the size of the ...
> ieeexplore.ieee.org/iel5/6993/18844/00869778.pdf -
> <u>Similar pages</u>
> [<u>More results from ieeexplore.ieee.org</u>]

<u>SkyMallSeasonal - Rolling **Tree Storage**</u>
Rolling **Tree Storage** from Improvements: Find Rolling **Tree Storage** at
SkyMall: Holiday ... However, if you find an **identical** new product
advertised for less, ...
www.skymall.com/shopping/detail.htm?pid=102190348&c=12020 - 92k -
<u>Cached</u> - <u>Similar pages</u>

[PDF] <u>B e s t - F i r s t F i x e d - D e p t h G a m e - T r e e Searc ...</u>
File Format: PDF/Adobe Acrobat - <u>View as HTML</u>
iments that compare them using **identical storage**. requirements. ....
Given a branching factor of w and a **tree** of depth d, the ...
dli.iiit.ac.in/ijcai/IJCAI-95-VOL%201/pdf/036.pdf - <u>Similar pages</u>

[PS] <u>Best-First Fixed-Depth Game-**Tree** Search in Practice</u>
File Format: Adobe PostScript - <u>View as Text</u>
identical storage. requirements. With dynamic move reordering, SSS*
is no longer ... game-**tree** searches are usually over a range of values.
...
www.cs.vu.nl/~aske/Papers/bfmms.ps.gz - <u>Similar pages</u>

<u>CRAFTS : Organizing & **Storage** : CD **Tree** House : DIY Network</u>
CD **Tree** House. Take your CD crate and measure it for a roof. Take two **identical** pieces of
Styrofoam and tape them together in the middle with a piece of ...
www.diynetwork.com/diy/cr_organizing_storage/
article/0,2025,DIY_13769_5088076,00.html - 38k - <u>Cached</u> - <u>Similar pages</u>

<u>Multiprocessor bridge having **storage** for spanning **tree** operation ...</u>

Multiprocessor bridge having **storage** for spanning **tree** operation mode ... that the contents
of all filtering databases within the bridge are kept **identical**. ...

www.patentstorm.us/patents/5761435-description.html - 37k - Cached - Similar pages

[PDF] Genetics of seed **storage** proteins in the love **tree** <Emphasis Type ...
File Format: PDF/Adobe Acrobat
Genetics of seed **storage** proteins in the love **tree** Cercis siliquastrum L. ... **identical** parts,
and a piece of cotyledon of about 2-3 mg was crushed ...
www.springerlink.com/index/U641167GR2410416.pdf - Similar pages

[PDF] Oracle7 Server
File Format: PDF/Adobe Acrobat - View as HTML
identical in every other respect. The purpose of this paper is to quantify the .... to realize
that with optimal **storage** the creation of b-**tree** indexes ...
www.oracle.com/technology/products/
bi/db/10g/pdf/twp_bi_dw_key_features_10gr2_0405.pdf - Similar pages

HDF Vset File Format for Persistent **Storage** of Semiconductor ...
Figure 5 HDF-Vset file hierarchy example for **Tree** Mesh ... the coordinates of the spatial
points not necessarily **identical** to their use in the Mesh VGroup. ...
www-tcad.stanford.edu/tcad/programs/hdfvset.html - 19k - Cached - Similar pages

1 2 3 4 5 6 7 8 9 10   **Next**

Try Google Desktop: search your computer as easily as you search the web.

| identical storage tree | Search |

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve

---

©2007 Google - Google Home - Advertising Programs - Business Solutions - About Google

P⟨⟩RTAL

USPTO

Search: ⊙ The ACM Digital Library   ○ The Guide

identical storage tree   **SEARCH**

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used: **identical storage tree**   Found **29,819** of **206,658**

| Sort results by | relevance ▼ | ● Save results to a Binder | Try an Advanced Search |
| Display results | expanded form ▼ | ⑦ Search Tips
☐ Open results in a new window | Try this search in The ACM Guide |

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   next
Best 200 shown   Relevance scale ☐▭▬◼◼

1 Physical storage structures: The K-D-B-tree: a search structure for large   ◼
multidimensional dynamic indexes
John T. Robinson
April 1981 **Proceedings of the 1981 ACM SIGMOD international conference on**
Management of data SIGMOD '81
Publisher: ACM Press
Full text available: 📄 pdf(723.91 KB)   Additional Information: full citation, abstract, references, citings

> The problem of retrieving multikey records via range queries from a large, dynamic index
> is considered. By *large* it is meant that most of the index must be stored on secondary
> memory. By *dynamic* it is meant that insertions and deletions are intermixed with queries,
> so that the index cannot be built beforehand. A new data structure, the *K-D-B-tree,* is
> presented as a solution to this problem. K-D-B-trees combine properties of K-D-trees and
> B-trees. It is expected that the mult ...

2 Physical storage structures: B+ trees and indexed sequential files: a performance   ◼
comparison
D. S. Batory
April 1981 **Proceedings of the 1981 ACM SIGMOD international conference on**
Management of data SIGMOD '81
Publisher: ACM Press
Full text available: 📄 pdf(858.82 KB)   Additional Information: full citation, abstract, references, citings

> An analytic method for comparing the performance of B+ trees and indexed sequential
> files is proposed. Preliminary results indicate that indexed sequential files may be more
> efficient than B+ trees in certain applications.

3 Decision Trees and Diagrams   ◼
Bernard M. E. Moret
December 1982 **ACM Computing Surveys (CSUR)**, Volume 14 Issue 4
Publisher: ACM Press
Full text available: 📄 pdf(2.68 MB)   Additional Information: full citation, references, citings, index terms

4 Sensor networks: Data storage placement in sensor networks   ◼
Bo Sheng, Qun Li, Weizhen Mao
May 2006 **Proceedings of the seventh ACM international symposium on Mobile ad**

hoc networking and computing MobiHoc '06
Publisher: ACM Press
Full text available: pdf(311.81 KB)   Additional Information: full citation, abstract, references, index terms

Data storage has become a important issue in sensor networks as a large amount of collected data need to be archived for future information retrieval. This paper introduces storage nodes to store the data collected from the sensors in their proximities. The storage nodes alleviate the heavy load of transmitting all the data to a central place for archiving and reduce the communicatio cost induced by the network query. This paper considers the storage node placement problem aiming to minimize the ...

Keywords: data query, data storage, wireless sensor networks

5 Self-assessment procedure XIII: a self-assessment procedure dealing with binary search trees and B-trees
Gopal K. Gupta
May 1984 **Communications of the ACM**, Volume 27 Issue 5
Publisher: ACM Press
Full text available: pdf(594.09 KB)   Additional Information: full citation, references, index terms

6 Region proximity in metric spaces and its use for approximate similarity search
Giuseppe Amato, Fausto Rabitti, Pasquale Savino, Pavel Zezula
April 2003 **ACM Transactions on Information Systems (TOIS)**, Volume 21 Issue 2
Publisher: ACM Press
Full text available: pdf(1.01 MB)   Additional Information: full citation, abstract, references, citings, index terms

Similarity search structures for metric data typically bound object partitions by ball regions. Since regions can overlap, a relevant issue is to estimate the proximity of regions in order to predict the number of objects in the regions' intersection. This paper analyzes the problem using a probabilistic approach and provides a solution that effectively computes the proximity through realistic heuristics that only require small amounts of auxiliary data. An extensive simulation to validate the t ...

Keywords: Approximation algorithms, approximate similarity search, metric data, metric trees, performance evaluation

7 Height-balanced trees of order (β, γ, δ)
Shou-Hsuan S. Huang
June 1985 **ACM Transactions on Database Systems (TODS)**, Volume 10 Issue 2
Publisher: ACM Press
Full text available: pdf(1.71 MB)   Additional Information: full citation, abstract, references, citings, index terms, review

We study restricted classes of B-trees, called H(&bgr;, &ggr;, &dgr;) trees. A class is defined by three parameters: &bgr;, the size of a node; &ggr;, the minimal number of grandsons a node must have; and &dgr;, the minimal number of leaves bottom nodes must have. This generalizes the brother condition of 2-3 brother trees in a uniform way to B-trees of higher order. The class of B-trees of order m is obtained by choosing &bgr; = m, &ggr; = (m/2)2 and &dgr; = m/2. An al ...

8 Implanting FFP trees in binary trees: An architectural proposal
Donald MacDavid Tolle

October 1981 **Proceedings of the 1981 conference on Functional programming languages and computer architecture FPCA '81**
Publisher: ACM Press
Full text available: 📄 pdf(611.57 KB)   Additional Information: full citation, abstract, references, index terms

The computer architecture described here was inspired by Magó's recently proposed cellular computer [1979] based on the Formal Functional Programming (FFP) languages introduced by Backus [1978]. Magó's machine is a binary tree of many simple processor/memory units, called cells. FFP programs are stored, one symbol per cell, in the leaf cells ("L" cells) of the machine. The machine fully exploits all the parallelism expressed in an FFP program, storage space permitt ...

9 Power proximity based key management for secure multicast in ad hoc networks ▄
Loukas Lazos, Radha Poovendran
January 2007 **Wireless Networks**, Volume 13 Issue 1
Publisher: Kluwer Academic Publishers
Full text available: 📄 pdf(875.28 KB)   Additional Information: full citation, abstract, references, index terms

As group-oriented services become the focal point of ad hoc network applications, securing the group communications becomes a default requirement. In this paper, we address the problem of group access in secure multicast communications for wireless ad hoc networks. We argue that energy expenditure is a scarce resource for the energy-limited ad hoc network devices and introduce a cross-layer approach for designing energy-efficient, balanced key distribution trees to perform key management. To con ...

Keywords: ad hoc networks, energy efficiency, key management, multicast, security

10 Improving duplicate elimination in storage systems ▄
Deepak R. Bobbarjung, Suresh Jagannathan, Cezary Dubnicki
November 2006 **ACM Transactions on Storage (TOS)**, Volume 2 Issue 4
Publisher: ACM Press
Full text available: 📄 pdf(481.90 KB)   Additional Information: full citation, abstract, references, index terms

Minimizing the amount of data that must be stored and managed is a key goal for any storage architecture that purports to be scalable. One way to achieve this goal is to avoid maintaining duplicate copies of the same data. Eliminating redundant data at the source by not writing data which has already been stored not only reduces storage overheads, but can also improve bandwidth utilization. For these reasons, in the face of today's exponentially growing data volumes, redundant data elimination t ...

Keywords: Rabin's fingerprints, Storage management, content-based addressing, duplicate elimination

11 Overlapping linear quadtrees: a spatio-temporal access method ▄
Theodoros Tzouramanis, Michael Vassilakopoulos, Yannis Manolopoulos
November 1998 **Proceedings of the 6th ACM international symposium on Advances in geographic information systems GIS '98**
Publisher: ACM Press
Full text available: 📄 pdf(837.72 KB)   Additional Information: full citation, references, citings, index terms

Keywords: B+ trees, quadtrees, access methods, indexing, linear quadtrees, overlapping, spatio-temporal databases, time/space performance, transaction time

12 Research sessions: compression: Dwarf: shrinking the PetaCube

Yannis Sismanis, Antonios Deligiannakis, Nick Roussopoulos, Yannis Kotidis

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**

Publisher: ACM Press

Full text available: pdf(1.38 MB)   Additional Information: full citation, abstract, references, citings, index terms

Dwarf is a highly compressed structure for computing, storing, and querying data cubes. Dwarf identifies prefix and suffix structural redundancies and factors them out by coalescing their store. Prefix redundancy is high on dense areas of cubes but suffix redundancy is significantly higher for sparse areas. Putting the two together fuses the exponential sizes of high dimensional full cubes into a dramatically condensed data structure. The elimination of suffix redundancy has an equally dramatic ...

13 New methods for dynamic storage allocation (Fast Fits)

C. J. Stephenson

October 1983 **ACM SIGOPS Operating Systems Review , Proceedings of the ninth ACM symposium on Operating systems principles SOSP '83**, Volume 17 Issue 5

Publisher: ACM Press

Full text available: pdf(192.01 KB)   Additional Information: full citation, abstract, citings, index terms

The classical methods for implementing dynamic storage allocation can be summarized thus First Fit and Best Fit The available blocks of storage are linked together in address order. Storage is allocated from the first available block of sufficient length, or from a block with the minimum excess length. Storage can be allocated or released in multiples of two words. In the long run, if numerous pieces of storage of more-or-less random lengths ar ...

14 Parallel implementation of an extended functional programming language on cellular tree and data flow architectures

Reza Sanati-Mehrizy, Krishna K. Agarwal, Afsaneh Minaie

March 1992 **Proceedings of the 1992 ACM/SIGAPP symposium on Applied computing: technological challenges of the 1990's SAC '92**

Publisher: ACM Press

Full text available: pdf(728.82 KB)   Additional Information: full citation, references, index terms

15 Distributed transactions for reliable systems

Alfred Z. Spector, Dean Daniels, Daniel Duchamp, Jeffrey L. Eppinger, Randy Pausch

December 1985 **ACM SIGOPS Operating Systems Review , Proceedings of the tenth ACM symposium on Operating systems principles SOSP '85**, Volume 19 Issue 5

Publisher: ACM Press

Full text available: pdf(1.44 MB)   Additional Information: full citation, references, citings, index terms

16 Table compression for tree automata

Jürgen Börstler, Ulrich Möncke, Reinhard Wilhelm

July 1991 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 13 Issue 3

Publisher: ACM Press

Full text available: pdf(1.15 MB)   Additional Information: full citation, references, index terms, review

Keywords: sparse tables, table compression, tree automata

17 Key addressable data storage

Roger Eggen
April 1998 **Proceedings of the 36th annual Southeast regional conference ACM-SE 36**
Publisher: ACM Press
Full text available: pdf(423.23 KB)   Additional Information: full citation, references, index terms

18 Packet classification on multiple fields

Pankaj Gupta, Nick McKeown
August 1999 **ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication SIGCOMM '99**, Volume 29 Issue 4
Publisher: ACM Press

Full text available: pdf(1.46 MB)   Additional Information: full citation, abstract, references, citings, index terms

Routers classify packets to determine which flow they belong to, and to decide what service they should receive. Classification may, in general, be based on an arbitrary number of fields in the packet header. Performing classification quickly on an arbitrary number of fields is known to be difficult, and has poor worst-case performance. In this paper, we consider a number of classifiers taken from real networks. We find that the classifiers contain considerable structure and redundancy that can ...

19 Industrial sessions: database internals - I: Query processing for SQL updates

César A. Galindo-Legaria, Stefano Stefani, Florian Waas
June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data SIGMOD '04**
Publisher: ACM Press
Full text available: pdf(73.48 KB)   Additional Information: full citation, abstract, references

A rich set of concepts and techniques has been developed in the context of query processing for the efficient and robust execution of queries. So far, this work has mostly focused on issues related to data-retrieval queries, with a strong backing on relational algebra. However, update operations can also exhibit a number of query processing issues, depending on the complexity of the operations and the volume of data to process. Such issues include lookup and matching of values, navigational vs. ...

20 Practical Aspects of Doubly Chained Trees for Retrieval

L. E. Stanfel
July 1972 **Journal of the ACM (JACM)**, Volume 19 Issue 3
Publisher: ACM Press
Full text available: pdf(614.61 KB)   Additional Information: full citation, references, index terms

Useful downloads: Adobe Acrobat   QuickTime   Windows Media Player   Real Player

Subscribe (Full Service)   Register (Limited Service, Free)   Login

Search: ◉ The ACM Digital Library   ○ The Guide

identical storage trees   **SEARCH**

THE ACM DIGITAL LIBRARY   📑 Feedback Report a problem Satisfaction
                                                                survey

Terms used: **identical storage trees**   Found **29,819** of **206,658**

Sort results   relevance   ▼   ❤ Save results to a Binder   Try an Advanced Search
by   Try this search in The ACM Guide
                                            ❓ Search Tips
Display   expanded form   ▼
results   ☐ Open results in a new
                                            window

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   next
Best 200 shown   Relevance scale ☐ ◩ ◪ ◼ ◼

1 Physical storage structures: The K-D-B-tree: a search structure for large   ◼
     multidimensional dynamic indexes
     John T. Robinson
     April 1981 **Proceedings of the 1981 ACM SIGMOD international conference on
                 Management of data SIGMOD '81**
     Publisher: ACM Press
     Full text available: 📄 pdf(723.91 KB)   Additional Information: full citation, abstract, references, citings

     The problem of retrieving multikey records via range queries from a large, dynamic index
     is considered. By *large* it is meant that most of the index must be stored on secondary
     memory. By *dynamic* it is meant that insertions and deletions are intermixed with queries,
     so that the index cannot be built beforehand. A new data structure, the *K-D-B-tree,* is
     presented as a solution to this problem. K-D-B-trees combine properties of K-D-trees and
     B-trees. It is expected that the mult ...

2 Physical storage structures: B+ trees and indexed sequential files: a performance   ◼
     comparison
     D. S. Batory
     April 1981 **Proceedings of the 1981 ACM SIGMOD international conference on
                 Management of data SIGMOD '81**
     Publisher: ACM Press
     Full text available: 📄 pdf(858.82 KB)   Additional Information: full citation, abstract, references, citings

     An analytic method for comparing the performance of B+ trees and indexed sequential
     files is proposed. Preliminary results indicate that indexed sequential files may be more
     efficient than B+ trees in certain applications.

3 Decision Trees and Diagrams   ◼
     Bernard M. E. Moret
     December 1982 **ACM Computing Surveys (CSUR)**, Volume 14 Issue 4
     Publisher: ACM Press
     Full text available: 📄 pdf(2.68 MB)   Additional Information: full citation, references, citings, index terms

4 Sensor networks: Data storage placement in sensor networks   ◼
     Bo Sheng, Qun Li, Weizhen Mao
     May 2006 **Proceedings of the seventh ACM international symposium on Mobile ad**

hoc networking and computing MobiHoc '06
Publisher: ACM Press
Full text available: pdf(311.81 KB)   Additional Information: full citation, abstract, references, index terms

Data storage has become a important issue in sensor networks as a large amount of collected data need to be archived for future information retrieval. This paper introduces storage nodes to store the data collected from the sensors in their proximities. The storage nodes alleviate the heavy load of transmitting all the data to a central place for archiving and reduce the communicatio cost induced by the network query. This paper considers the storage node placement problem aiming to minimize the ...

Keywords: data query, data storage, wireless sensor networks

5 Self-assessment procedure XIII: a self-assessment procedure dealing with binary search trees and B-trees
Gopal K. Gupta
May 1984 **Communications of the ACM**, Volume 27 Issue 5
Publisher: ACM Press
Full text available: pdf(594.09 KB)   Additional Information: full citation, references, index terms

6 Region proximity in metric spaces and its use for approximate similarity search
Giuseppe Amato, Fausto Rabitti, Pasquale Savino, Pavel Zezula
April 2003 **ACM Transactions on Information Systems (TOIS)**, Volume 21 Issue 2
Publisher: ACM Press
Full text available: pdf(1.01 MB)   Additional Information: full citation, abstract, references, citings, index terms

Similarity search structures for metric data typically bound object partitions by ball regions. Since regions can overlap, a relevant issue is to estimate the proximity of regions in order to predict the number of objects in the regions' intersection. This paper analyzes the problem using a probabilistic approach and provides a solution that effectively computes the proximity through realistic heuristics that only require small amounts of auxiliary data. An extensive simulation to validate the t ...

Keywords: Approximation algorithms, approximate similarity search, metric data, metric trees, performance evaluation

7 Height-balanced trees of order (β, γ, δ)
Shou-Hsuan S. Huang
June 1985 **ACM Transactions on Database Systems (TODS)**, Volume 10 Issue 2
Publisher: ACM Press
Full text available: pdf(1.71 MB)   Additional Information: full citation, abstract, references, citings, index terms, review

We study restricted classes of B-trees, called H(&bgr;, &ggr;, &dgr;) trees. A class is defined by three parameters: &bgr;, the size of a node; &ggr;, the minimal number of grandsons a node must have; and &dgr;, the minimal number of leaves bottom nodes must have. This generalizes the brother condition of 2-3 brother trees in a uniform way to B-trees of higher order. The class of B-trees of order m is obtained by choosing &bgr; = m, &ggr; = (m/2)2 and &dgr; = m/2. An al ...

8 Implanting FFP trees in binary trees: An architectural proposal
Donald MacDavid Tolle

October 1981 **Proceedings of the 1981 conference on Functional programming languages and computer architecture FPCA '81**
Publisher: ACM Press
Full text available: pdf(611.57 KB)   Additional Information: full citation, abstract, references, index terms

The computer architecture described here was inspired by Magó's recently proposed cellular computer [1979] based on the Formal Functional Programming (FFP) languages introduced by Backus [1978]. Magó's machine is a binary tree of many simple processor/memory units, called cells. FFP programs are stored, one symbol per cell, in the leaf cells ("L" cells) of the machine. The machine fully exploits all the parallelism expressed in an FFP program, storage space permitt ...

9   Power proximity based key management for secure multicast in ad hoc networks
Loukas Lazos, Radha Poovendran
January 2007 **Wireless Networks**, Volume 13 Issue 1
Publisher: Kluwer Academic Publishers
Full text available: pdf(875.28 KB)   Additional Information: full citation, abstract, references, index terms

As group-oriented services become the focal point of ad hoc network applications, securing the group communications becomes a default requirement. In this paper, we address the problem of group access in secure multicast communications for wireless ad hoc networks. We argue that energy expenditure is a scarce resource for the energy-limited ad hoc network devices and introduce a cross-layer approach for designing energy-efficient, balanced key distribution trees to perform key management. To con ...

Keywords: ad hoc networks, energy efficiency, key management, multicast, security

10   Improving duplicate elimination in storage systems
Deepak R. Bobbarjung, Suresh Jagannathan, Cezary Dubnicki
November 2006 **ACM Transactions on Storage (TOS)**, Volume 2 Issue 4
Publisher: ACM Press
Full text available: pdf(481.90 KB)   Additional Information: full citation, abstract, references, index terms

Minimizing the amount of data that must be stored and managed is a key goal for any storage architecture that purports to be scalable. One way to achieve this goal is to avoid maintaining duplicate copies of the same data. Eliminating redundant data at the source by not writing data which has already been stored not only reduces storage overheads, but can also improve bandwidth utilization. For these reasons, in the face of today's exponentially growing data volumes, redundant data elimination t ...

Keywords: Rabin's fingerprints, Storage management, content-based addressing, duplicate elimination

11   Overlapping linear quadtrees: a spatio-temporal access method
Theodoros Tzouramanis, Michael Vassilakopoulos, Yannis Manolopoulos
November 1998 **Proceedings of the 6th ACM international symposium on Advances in geographic information systems GIS '98**
Publisher: ACM Press
Full text available: pdf(837.72 KB)   Additional Information: full citation, references, citings, index terms

Keywords: B+ trees, quadtrees, access methods, indexing, linear quadtrees, overlapping, spatio-temporal databases, time/space performance, transaction time

12 Research sessions: compression: Dwarf: shrinking the PetaCube ■

Yannis Sismanis, Antonios Deligiannakis, Nick Roussopoulos, Yannis Kotidis

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**

Publisher: ACM Press

Full text available: 📄 pdf(1.38 MB)   Additional Information: full citation, abstract, references, citings, index terms

Dwarf is a highly compressed structure for computing, storing, and querying data cubes. Dwarf identifies prefix and suffix structural redundancies and factors them out by coalescing their store. Prefix redundancy is high on dense areas of cubes but suffix redundancy is significantly higher for sparse areas. Putting the two together fuses the exponential sizes of high dimensional full cubes into a dramatically condensed data structure. The elimination of suffix redundancy has an equally dramatic ...

13 New methods for dynamic storage allocation (Fast Fits) ■

C. J. Stephenson

October 1983 **ACM SIGOPS Operating Systems Review , Proceedings of the ninth ACM symposium on Operating systems principles SOSP '83**, Volume 17 Issue 5

Publisher: ACM Press

Full text available: 📄 pdf(192.01 KB)   Additional Information: full citation, abstract, citings, index terms

The classical methods for implementing dynamic storage allocation can be summarized thus First Fit and Best Fit The available blocks of storage are linked together in address order. Storage is allocated from the first available block of sufficient length, or from a block with the minimum excess length. Storage can be allocated or released in multiples of two words. In the long run, if numerous pieces of storage of more-or-less random lengths ar ...

14 Parallel implementation of an extended functional programming language on cellular ■ tree and data flow architectures

Reza Sanati-Mehrizy, Krishna K. Agarwal, Afsaneh Minaie

March 1992 **Proceedings of the 1992 ACM/SIGAPP symposium on Applied computing: technological challenges of the 1990's SAC '92**

Publisher: ACM Press

Full text available: 📄 pdf(728.82 KB)   Additional Information: full citation, references, index terms

15 Distributed transactions for reliable systems ■

Alfred Z. Spector, Dean Daniels, Daniel Duchamp, Jeffrey L. Eppinger, Randy Pausch

December 1985 **ACM SIGOPS Operating Systems Review , Proceedings of the tenth ACM symposium on Operating systems principles SOSP '85**, Volume 19 Issue 5

Publisher: ACM Press

Full text available: 📄 pdf(1.44 MB)   Additional Information: full citation, references, citings, index terms

16 Table compression for tree automata ■

Jürgen Börstler, Ulrich Möncke, Reinhard Wilhelm

July 1991 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 13 Issue 3

Publisher: ACM Press

Full text available: 📄 pdf(1.15 MB)   Additional Information: full citation, references, index terms, review

Keywords: sparse tables, table compression, tree automata

17 Key addressable data storage

Roger Eggen

April 1998 **Proceedings of the 36th annual Southeast regional conference ACM-SE 36**

Publisher: ACM Press

Full text available: pdf(423.23 KB)   Additional Information: full citation, references, index terms

18 Packet classification on multiple fields

Pankaj Gupta, Nick McKeown

August 1999 **ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication SIGCOMM '99**, Volume 29 Issue 4

Publisher: ACM Press

Full text available: pdf(1.46 MB)   Additional Information: full citation, abstract, references, citings, index terms

Routers classify packets to determine which flow they belong to, and to decide what service they should receive. Classification may, in general, be based on an arbitrary number of fields in the packet header. Performing classification quickly on an arbitrary number of fields is known to be difficult, and has poor worst-case performance. In this paper, we consider a number of classifiers taken from real networks. We find that the classifiers contain considerable structure and redundancy that can ...

19 Industrial sessions: database internals - I: Query processing for SQL updates

César A. Galindo-Legaria, Stefano Stefani, Florian Waas

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data SIGMOD '04**

Publisher: ACM Press

Full text available: pdf(73.48 KB)   Additional Information: full citation, abstract, references

A rich set of concepts and techniques has been developed in the context of query processing for the efficient and robust execution of queries. So far, this work has mostly focused on issues related to data-retrieval queries, with a strong backing on relational algebra. However, update operations can also exhibit a number of query processing issues, depending on the complexity of the operations and the volume of data to process. Such issues include lookup and matching of values, navigational vs. ...

20 Practical Aspects of Doubly Chained Trees for Retrieval

L. E. Stanfel

July 1972 **Journal of the ACM (JACM)**, Volume 19 Issue 3

Publisher: ACM Press

Full text available: pdf(614.61 KB)   Additional Information: full citation, references, index terms

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10 next

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

P⊘RTAL

USPTO

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used: **migratable** **storage** **root** **key**   Found **36,313** of **206,658**

Sort results by   [relevance ▼]   ● Save results to a Binder   Try an Advanced Search
Display results   [expanded form ▼]   ? Search Tips   Try this search in The ACM Guide
                                          ☐ Open results in a new window

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   next
Best 200 shown   Relevance scale ☐▨▨■■

1 Applications and compliance: Virtual monotonic counters and count-limited objects   ■
    using a TPM without a trusted OS
    Luis F. G. Sarmenta, Marten van Dijk, Charles W. O'Donnell, Jonathan Rhodes, Srinivas
    Devadas
    November 2006 **Proceedings of the first ACM workshop on Scalable trusted computing
               STC '06**
    Publisher: ACM Press
    Full text available: 🔲 pdf(447.59 KB)   Additional Information: full citation, abstract, references, index terms

    A trusted monotonic counter is a valuable primitive that enables a wide variety of highly
    scalable offline and decentralized applications that would otherwise be prone to replay
    attacks, including offline payment, e-wallets, virtual trusted storage, and digital rights
    management (DRM). In this paper, we show how one can implement a very large number
    of *virtual* monotonic counters on an untrusted machine with a Trusted Platform Module
    (TPM) or similar device, without relying on a trusted OS ...

    Keywords: certified execution, e-wallet memory integrity checking, key delegation,
    stored-value, trusted storage

2 Computer forensics (CF): The advent of trusted computing: implications for digital   ■
    forensics
    Mike Burmester, Judie Mulholland
    April 2006 **Proceedings of the 2006 ACM symposium on Applied computing SAC '06**
    Publisher: ACM Press
    Full text available: 🔲 pdf(137.02 KB)   Additional Information: full citation, abstract, references, index terms

    The release of computer hardware devices based on "trusted computing" technologies is
    heralding a paradigm shift that will have profound implications for digital forensics. In this
    paper, we map out the contours of a trusted environment in order to establish the context
    for the paper. This is followed by the main components of the TC architecture with an
    emphasis on the Trusted Platform and the Trusted Platform Module (TPM). The next
    section presents a synopsis based on three threat models, *v ...*

    Keywords: *cybercrime, data recovery, encryption, file systems, forensics, specifications,
    trusted computing*

3 Access management for distributed systems: Peer-to-peer access control architecture using trusted computing technology
Ravi Sandhu, Xinwen Zhang
June 2005 **Proceedings of the tenth ACM symposium on Access control models and technologies SACMAT '05**
Publisher: ACM Press

Full text available: pdf(215.48 KB)   Additional Information: full citation, abstract, references, citings, index terms, review

It has been recognized for some time that software alone does not provide an adequate foundation for building a high-assurance trusted platform. The emergence of industry-standard trusted computing technologies promises a revolution in this respect by providing roots of trust upon which secure applications can be developed. These technologies offer a particularly attractive platform for security in peer-to-peer environments. In this paper we propose a trusted computing architecture to enforce ac ...

Keywords: access control, policy enforcement, security architecture, trusted computing

4 Secure systems: Energy and execution time analysis of a software-based trusted platform module
Najwa Aaraj, Anand Raghunathan, Srivaths Ravi, Niraj K. Jha
April 2007 **Proceedings of the conference on Design, automation and test in Europe DATE '07**
Publisher: ACM Press

Full text available: pdf(838.82 KB)   Additional Information: full citation, abstract, references

Trusted platforms have been proposed as a promising approach to enhance the security of general-purpose computing systems. However, for many resource-constrained embedded systems, the size and cost overheads of a separate Trusted Platform Module (TPM) chip are not acceptable. One alternative is to use a software-based TPM (SW-TPM), which implements TPM functions using software that executes in a protected execution domain on the embedded processor itself. However, since many embedded systems ...

5 Architechture: Towards an open, trusted digital rights management platform
Andrew Cooper, Andrew Martin
October 2006 **Proceedings of the ACM workshop on Digital rights management DRM '06**
Publisher: ACM Press

Full text available: pdf(417.51 KB)   Additional Information: full citation, abstract, references, index terms

Trusted computing has received criticism from those who fear it will be used by influential market forces to exert power over the software used on consumer platforms. This paper describes an open architecture for digital rights management (DRM) enforcement on trusted computing platforms that empowers the consumer to select their operating-system and applications, including open-source options, without weakening the strength of the security functions. A key component in the architecture is a secu ...

Keywords: DRM, digital rights management, mandatory access controls, trusted computing, virtual machines

6 Applications and compliance: TCG inside?: a note on TPM specification compliance
Ahmad-Reza Sadeghi, Marcel Selhorst, Christian Stüble, Christian Wachsmann, Marcel Winandy, Horst Görtz
November 2006 **Proceedings of the first ACM workshop on Scalable trusted computing**

STC '06
Publisher: ACM Press
Full text available: 📄 pdf(587.22 KB)   Additional Information: full citation, abstract, references, index terms

The Trusted Computing Group (TCG) has addressed a new generation of computing platforms employing both supplemental hardware and software with the primary goal to improve the security and the trustworthiness of future IT systems. The core component of the TCG proposal is the Trusted Platform Module (TPM) providing certain cryptographic functions. Many vendors currently equip their platforms with a TPM claiming to be TCG compliant. However, there is no feasible way for application developers and ...

Keywords: TPM, compliance, test, trusted computing

7   Improving the browsing experience: WebPod: persistent Web browsing sessions with   ▭
     pocketable storage devices
     Shaya Potter, Jason Nieh
     May 2005 **Proceedings of the 14th international conference on World Wide Web
            WWW '05**
     Publisher: ACM Press
     Full text available: 📄 pdf(166.59 KB)   Additional Information: full citation, abstract, references, index terms

We present WebPod, a portable system that enables mobile users to use the same persistent, personalized web browsing session on any Internet-enabled device. No matter what computer is being used, WebPod provides a consistent browsing session, maintaining all of a user's plugins, bookmarks, browser web content, open browser windows, and browser configuration options and preferences. This is achieved by leveraging rapid improvements in capacity, cost, and size of portable storage devices. WebPod p ...

Keywords: checkpoint/restart, portable storage, process migration, virtualization, web browsing

8   System support for pervasive applications   ▭
     Robert Grimm, Janet Davis, Eric Lemar, Adam Macbeth, Steven Swanson, Thomas Anderson,
     Brian Bershad, Gaetano Borriello, Steven Gribble, David Wetherall
     November 2004 **ACM Transactions on Computer Systems (TOCS)**, Volume 22 Issue 4
     Publisher: ACM Press
     Full text available: 📄 pdf(1.82 MB)   Additional Information: full citation, abstract, references, citings, index terms

Pervasive computing provides an attractive vision for the future of computing. Computational power will be available everywhere. Mobile and stationary devices will dynamically connect and coordinate to seamlessly help people in accomplishing their tasks. For this vision to become a reality, developers must build applications that constantly adapt to a highly dynamic computing environment. To make the developers' task feasible, we present a system architecture for pervasive computing, called & ...

Keywords: Asynchronous events, checkpointing, discovery, logic/operation pattern, migration, one.world, pervasive computing, structured I/O, tuples, ubiquitous computing

9   Migration: The design and implementation of Zap: a system for migrating computing   ▭
     environments
     Steven Osman, Dinesh Subhraveti, Gong Su, Jason Nieh
     December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Publisher: ACM Press

Full text available: pdf(2.06 MB)   Additional Information: full citation, abstract, references, cited by, index terms

> We have created Zap, a novel system for transparent migration of legacy and networked applications. Zap provides a thin virtualization layer on top of the operating system that introduces pods, which are groups of processes that are provided a consistent, virtualized view of the system. This decouples processes in pods from dependencies to the host operating system and other processes on the system. By integrating Zap virtualization with a checkpoint-restart mechanism, Zap can migrate a pod of p ...

10 Caching in networks (extended abstract)
Friedhelm Meyer auf der Heide, Berthold Vöcking, Matthias Westermann
February 2000 **Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms SODA '00**
Publisher: Society for Industrial and Applied Mathematics
Full text available: pdf(1.09 MB)   Additional Information: full citation, references, citings, index terms

11 Session 7: OS architecture II: Easing the management of data-parallel systems via adaptation
David Petrou, Khalil Amiri, Gregory R. Ganger, Garth A. Gibson
September 2000 **Proceedings of the 9th workshop on ACM SIGOPS European workshop: beyond the PC: new challenges for the operating system EW 9**
Publisher: ACM Press
Full text available: pdf(146.51 KB)   Additional Information: full citation, abstract, references

> In recent years we have seen an enormous growth in the size and prevalence of data processing workloads [Fayyad 1998, Gray 1997]. The picture that is becoming increasingly common is depicted in Figure 1. In it, organizations or resourceful individuals provide services via a set of loosely-coupled workstation nodes. The service is usually some form of data-mining like searching, filtering, or image recognition. Clients, which could be machines running web browsers, not only initiate requests, but ...

12 Attestation and binding: A protocol for property-based attestation
Liqun Chen, Rainer Landfermann, Hans Löhr, Markus Rohe, Ahmad-Reza Sadeghi, Christian Stüble, Horst Görtz
November 2006 **Proceedings of the first ACM workshop on Scalable trusted computing STC '06**
Publisher: ACM Press
Full text available: pdf(279.99 KB)   Additional Information: full citation, abstract, references, index terms

> The Trusted Computing Group (TCG) has issued several specifications to enhance the architecture of common computing platforms by means of new functionalities, amongst others the (binary) attestation to verify the integrity of a (remote) computing platform/application. However, as pointed out recently, the binary attestation has some shortcomings, in particular when used for applications: First, it reveals information about the configuration of a platform (hardware and software) or application. T ...

> **Keywords**: TCG binary attestation, property-based attestation, security kernels, zero-knowledge proof of knowledge

13 A scalable key-management scheme with minimizing key storage for secure group communications

Yuh-Min Tseng
November 2003 **International Journal of Network Management**, Volume 13 Issue 6
Publisher: John Wiley & Sons, Inc.
Full text available: pdf(104.15 KB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>index terms</u>

> Recently, many group communication services have become the focus for future
> developments in the Internet and wireless network applications, such as video-
> conferencing, collaborative work, networking games or online videos. In particular, these
> applications require data delivery from one sender to a large number of authorized
> receivers. Therefore, secure multicast communication will become an important
> networking issue in the future. Using a common encryption key only known by authorized
> members ...

14 <u>Zones, contracts and absorbing changes: an approach to software evolution</u>

Huw Evans, Peter Dickman
October 1999 **ACM SIGPLAN Notices , Proceedings of the 14th ACM SIGPLAN
conference on Object-oriented programming, systems, languages, and
applications OOPSLA '99**, Volume 34 Issue 10
Publisher: ACM Press

Full text available: pdf(2.46 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index
terms</u>

> This paper describes a novel approach to managing the evolution of distributed, persistent
> systems at run-time. This is achieved by partitioning a system into disjoint zones, each of
> which can be evolved without affecting code in any other. Contracts are defined between
> zones, making type-level interdependencies and inter-zone communication explicit.
> Programmer supplied code is added to the running system, at the boundary between
> zones, to constrain the sco ...

15 <u>Supporting dynamic data structures on distributed-memory machines</u>

Anne Rogers, Martin C. Carlisle, John H. Reppy, Laurie J. Hendren
March 1995 **ACM Transactions on Programming Languages and Systems (TOPLAS)**,
Volume 17 Issue 2
Publisher: ACM Press

Full text available: pdf(2.05 MB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index
terms</u>, <u>review</u>

> Compiling for distributed-memory machines has been a very active research area in
> recent years. Much of this work has concentrated on programs that use arrays as their
> primary data structures. To date, little work has been done to address the problem of
> supporting programs that use pointer-based dynamic data structures. The techniques
> developed for supporting SPMD execution of array-based programs rely on the fact that
> arrays are statically defined and directly addressable. Recursive data s ...

> **Keywords**: dynamic data structures

16 <u>Short papers -- storage survivability: Toward securing untrusted storage without
public-key operations</u>

Dalit Naor, Amir Shenhav, Avishai Wool
November 2005 **Proceedings of the 2005 ACM workshop on Storage security and
survivability StorageSS '05**
Publisher: ACM Press
Full text available: pdf(344.77 KB)   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>index terms</u>

> Adding security capabilities to shared, remote and untrusted storage file systems leads to
> performance degradation that limits their use. Public-key cryptographic primitives, widely
> used in such file systems, are known to have worse performance than their symmetric

key counterparts. In this paper we examine design alternatives that avoid public-key cryptography operations to achieve better performance. We present the trade-offs and limitations that are introduced by these substitutions.

Keywords: network attached storage, secure file systems

17 Decentralized storage systems: Farsite: federated, available, and reliable storage for an incompletely trusted environment
Atul Adya, William J. Bolosky, Miguel Castro, Gerald Cermak, Ronnie Chaiken, John R. Douceur, Jon Howell, Jacob R. Lorch, Marvin Theimer, Roger P. Wattenhofer
December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI
Publisher: ACM Press

Full text available: pdf(1.87 MB)   Additional Information: full citation, abstract, references, cited by, index terms

Farsite is a secure, scalable file system that logically functions as a centralized file server but is physically distributed among a set of untrusted computers. Farsite provides file availability and reliability through randomized replicated storage; it ensures the secrecy of file contents with cryptographic techniques; it maintains the integrity of file and directory data with a Byzantine-fault-tolerant protocol; it is designed to be scalable by using a distributed hint mechanism and delegatio ...

18 Sensornet services: TSAR: a two tier sensor storage architecture using interval skip graphs
Peter Desnoyers, Deepak Ganesan, Prashant Shenoy
November 2005 **Proceedings of the 3rd international conference on Embedded networked sensor systems SenSys '05**
Publisher: ACM Press

Full text available: pdf(444.47 KB)   Additional Information: full citation, abstract, references, citings, index terms

Archival storage of sensor data is necessary for applications that query, mine, and analyze such data for interesting features and trends. We argue that existing storage systems are designed primarily for flat hierarchies of homogeneous sensor nodes and do not fully exploit the multi-tier nature of emerging sensor networks, where an application can comprise tens of tethered proxies, each managing tens to hundreds of untethered sensors. We present *TSAR,* a fundamentally different storage ar ...

Keywords: archival storage, indexing methods, wireless sensor networks

19 Modeling the storage architectures of commercial database systems
D. S. Batory
December 1985 **ACM Transactions on Database Systems (TODS)**, Volume 10 Issue 4
Publisher: ACM Press

Full text available: pdf(4.46 MB)   Additional Information: full citation, abstract, references, citings, index terms, review

Modeling the storage structures of a DBMS is a prerequisite to understanding and optimizing database performance. Previously, such modeling was very difficult because the fundamental role of conceptual-to-internal mappings in DBMS implementations went unrecognized. In this paper we present a model of physical databases, called the transformation model, that makes conceptual-to-internal mappings explicit. By exposing such mappings, we show that it is possible to model the storage ...

20

A high performance, universal, key associative access method

David B. Lomet

May 1983 **ACM SIGMOD Record , Proceedings of the 1983 ACM SIGMOD international conference on Management of data SIGMOD '83**, Volume 13 Issue 4

Publisher: ACM Press

Full text available: pdf(1.81 MB)   Additional Information: full citation, abstract, references, citings

A new file organization is proposed that combines the advantages of digital B-trees and extendible hashing methods into one organization that can be used universally. The method, like these predecessors, relies on digital searching. The key notions are: (i) that multipage nodes are addressed by the root and can have both data and index entries, the mix of entries changing over time; and (ii) that these nodes can be doubled with file growth and, when this occurs, data nodes at the next level of t ...

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   next

Useful downloads: Adobe Acrobat   QuickTime   Windows Media Player   Real Player